

REMARKS

Pending Claims:

Claims 1-5, 7, 9, 10, 13, and 16-24 are pending in this application as amended. Claims 4, 10, 13, and 16-18 are currently amended. Claims 1, 2, 7, 9, and 19-24 were previously presented. Claims 3 and 5 are original. Claims 6, 8, 11, 12, 14, and 15 were canceled previously.

Rejections under 35 U.S.C. §103(a):

Claims Pertaining to Expansion

Claims 1, 2, 9, and 19-24 were rejected as being obvious over the combination of Lubbers et al. (US 6,880,052) and Bridge (US 6,530,035). Applicant respectfully disagrees with and traverses these rejections.

With respect to independent claim 1 (for a program storage device), the Examiner cites Lubbers as teaching “wherein a LUN may be resized and the resizing is automatically propagated to other members of the copy set.” Office Action (OA) p. 3. Bridges is cited as teaching respective sequences for expanding and shrinking logical volumes. This combination misses critical limitations of Applicant’s invention. Notice the order that Lubbers teaches. First the LUN is resized. Then the resizing is automatically propagated to other members of the copy set. This is radically different from what Applicant does.

Suppose a first virtual disk and a second virtual disk mirror each other, and suppose that write commands are coming in at random intervals to the mirrored pair from a host. And first suppose we are dealing with a capacity expansion of the mirrored pair. Suppose, as Lubbers suggests, the first virtual disk is resized first. Now suppose that a write request comes in, and the first virtual disk uses the newly added space to satisfy this write request. Now the two virtual disks are out of

synchronization. Just because the resizing is “automatically propagated” to the second virtual disk, there is no guarantee that the resizing is propagated synchronously to second virtual disk. Automatically might mean that the changes are propagated later, or even that both disks are automatically quiesced to catch the second up with the first. Lubbers never explains how this automatic resizing of copies in the copy set is achieved. However, the one discussion (col. 9, line 61 – col. 10, line 21) that seems relevant states, “Asynchronous operation refers to an operation mode in which a modification to one member of a copy set can be propagated to other members of the copy set.” In other words, you can follow Lubbers and use Bridge’s approach to expand, and still wind up with an automated approach that is asynchronous.

In Applicant’s approach, the space is made available for both virtual disks before the new size is reported. The reallocation goes on behind the scenes. When the new size is reported, both mirroring virtual disks are available to process write requests, without any loss of synchronization. This is a major difference, which does not follow from either reference or their combination. Note that Bridge deals with resizing individual logical volumes, and does not specify the order of how a pair of mirroring disks should be handled. You can combine Lubbers and Bridge and get a scheme that is asynchronous, and not even close to Applicant’s invention.

Independent claims 20 (for a method) and 23 (for an apparatus) are analogous to claim 1, with respect to the order in which expansion of the mirrored virtual disks is executed. Thus, claims 1, 20, 23, and all the dependent claims (2, 9, 19, 21, 22, and 24) that depend on them should be allowable.

Claims 3, 17, and 18 were rejected as obvious over the combination of Lubbers, Bridge, and Cabrera (US 6,629,202). With respect to claim 3, Cabrera was cited by the Examiner as teaching “specifying a size for a virtual disk and mapping the size of the virtual disk is performed by an operating system.” OA par. 13. The Cabrera reference is irrelevant to the above argument, and so claim 3 should be allowed.

Claim 17 should be allowed for a reason independent of claim 1. The method in claim 17 includes the step of “providing by the source virtual disk continuous availability for normal disk access operations between the step of receiving a request and the step of reporting the new storage size of the source virtual disk.” The Examiner states that “Lubbers discloses the host can continuously write to source (col. 12, line 38-col. 13, line 15).” This is irrelevant to claim 17 for two reasons:

- Figure 6, to which the cited language refers, pertains to creation of a new destination virtual disk (items 607, 609, 611, 613, 615, 617, 619), and not to resizing of a mirrored pair or set; and
- The process which Figure 6 describes includes quiescing the source virtual disk (item 617) and copying data in background (item 619) – this is anything but “continuous availability.”

Cabrera is referenced as disclosing “logical volumes and their plex are dynamically mapped and resized under the control of the operating system without system disruption.” Cabrera does not disclose the relative timing or resizing and reporting in a mirrored pair required by claim 1. Consequently, claims 17 and 18 should be allowed independently of claim 1.

Claim 18 should be allowable independently of claims 1 and 17. Claim 18 provides for continuous mirroring between the step of receiving a request and the step of reporting the new size of the source virtual disk. The portions of Lubbers and Cabrera cited by the Examiner as teaching this do not pertain to a resizing operation.

Claims Pertaining to Shrinking

Independent claims 4, 7, 10, 13, and 16 have been amended to be consistent with the specification, for example, par. [0049], which states that “The process is reversed for dynamically shrinking mirrored virtual disks”

Claim 4 was rejected as obvious over the combination of Lubbers, Bridge, and DeKoning (US 6,275,898). DeKoning was cited in reference to aspects of claim 4 dealing with RAID storage systems. Otherwise, the grounds for rejection parallel those of claim 1. Specifically, Lubbers is cited as disclosing “resizing members of a copy set dynamically wherein any change made to one LUN member of a copy set is automatically propagated to the other members.” Lubbers states, “when changes are made to a dependent attribute [such as size] of one member of a copy set, the change is made automatically to each other member of the copy set.” Again, the term “automatically” is ambiguous, and might simply mean under computer control. Lubbers does not specify the order for reporting a reduction in the size of the copy set. Bridge deals with shrinking a single logical volume, and so does not provide this critical missing piece. Applicant’s disclosure is specific: “The process is reversed for dynamically shrinking mirrored virtual disks in a RAID storage system, with the exception that when downsizing, you may need to shrink beyond the granularity that you expanded by.” Par. [0049].

In summary, key to claim 4 is that these are mirrored virtual disks. When they are downsized, the downsizing must be reported in a specific order relative to the actual shrinking of the available storage for the virtual disks. Bridge describes a process for downloading individual logical volumes. Lubbers is silent about the timing of changing the reported size of the virtual disks and shrinking individual volumes in a copy set. The order is important. By immediately reporting the virtual disks in the mirrored pair as downsized, the storage system can continue without interruption in service. Whatever reductions in number or sizes of physical disks used to implement the virtualization scheme can be done afterwards.

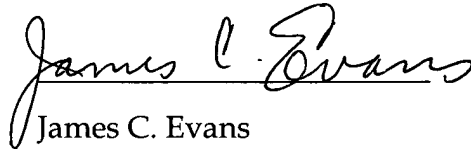
The above rationale for allowance applies to independent claims 7, 10, 13, and 16 as well. Dependent claim 5, which depends upon claim 4, should also be allowable.

CONCLUSION

All of the claims remaining in this application should now be seen to be in condition for allowance. The prompt issuance of a notice to that effect is solicited.

Date: 2/3/2010

Respectfully submitted,



James C. Evans
Registration No. 56,730
Beck & Tysver, P.L.L.C.
2900 Thomas Ave., #100
Minneapolis, MN 55416
Telephone: (612) 915-7006
Fax: (612) 915-9637